

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
3 April 2003 (03.04.2003)

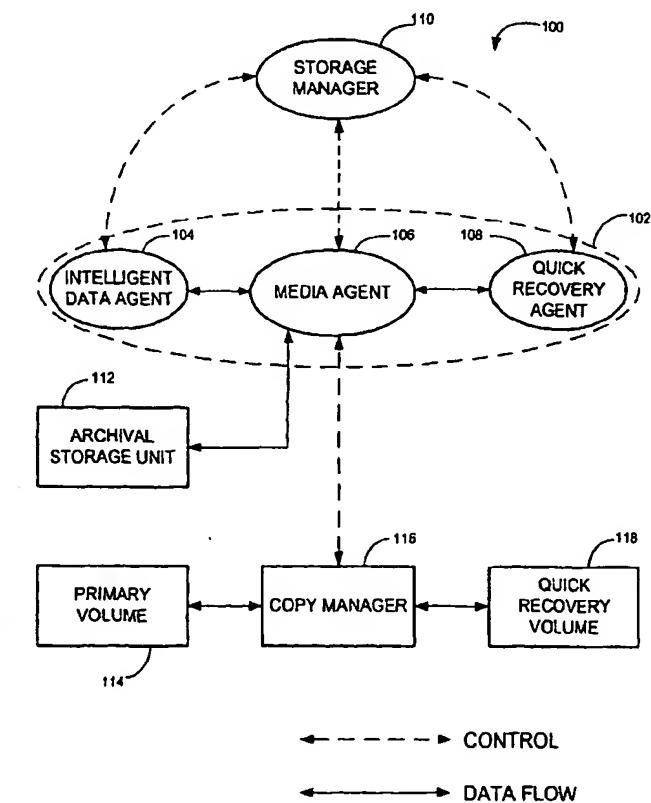
PCT

(10) International Publication Number
WO 03/028183 A1

- (51) International Patent Classification⁷: H02H 3/05, G06F 12/00
- (21) International Application Number: PCT/US02/31206
- (22) International Filing Date: 30 September 2002 (30.09.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 60/326,021 28 September 2001 (28.09.2001) US
- (71) Applicant: COMMVAULT SYSTEMS, INC. [US/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US).
- (72) Inventors: PRAHLAD, Anand; 3 Bucknell Drive, East Brunswick, NJ 08816 (US). PITTAUGA, Ivan;
- (74) Agent: OSTROW, Seth, H.; Brown Raysman Millstein Felder & Steiner, LLP, 20th Floor, 900 Third Avenue, New York, NY 10022 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, IIR, IIU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SI, SZ, TZ, UG, ZM, ZW),

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR GENERATING AND MANAGING QUICK RECOVERY VOLUMES



(57) Abstract: The invention relates to computer readable medium storing program code which when executed on a computer causes the computer to perform a method for creating a quick recovery volume (118) of a primary data set used by a first computer in a backup storage system, which includes identifying a snapshot image of the primary data set generated by a snapshot application, creating the quick recovery volume (118) of the primary data set from the snapshot image of the primary data set (114) and controlling transfer of data from the first computer to an archival storage unit (112). In one embodiment, the invention provides a method for creating a quick recovery volume (118) of a primary data set that includes creating a snapshot image of the primary data set (114) and creating a quick recovery volume (118) of the primary data set from the snapshot image of the primary data set (114). In another embodiment, the invention provides a user interface screen enabling browser style browsing and recovery of quick recovery volumes (118) and snapshot images.

WO 03/028183 A1



Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

Published:

— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

**SYSTEM AND METHOD FOR GENERATING AND MANAGING QUICK
RECOVERY VOLUMES**

PRIORITY CLAIM

5 This application claims priority from United States Provisional Application No. 60/326,021, entitled "METHOD FOR MANAGING SNAPSHOTS GENERATED BY AN OPERATING SYSTEM OR OTHER APPLICATION", filed September 28, 2001. The entire contents of the Provisional Application 60/326,021 are hereby incorporated herein by reference.

10 **COPYRIGHT NOTICE**

A portion of the disclosure of this patent document contains material, which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all copyright rights whatsoever.

RELATED APPLICATIONS

This application is related to the following pending applications:

- Application Serial No. 09/610,738, titled MODULAR BACKUP AND RETRIEVAL SYSTEM USED IN CONJUNCTION WITH A STORAGE AREA NETWORK, filed July 6, 2000, attorney docket number 044463-002;
- Application Serial No. 09/609,977, titled MODULAR BACKUP AND RETRIEVAL SYSTEM WITH AN INTEGRATED STORAGE AREA FILING SYSTEM, filed August 5, 2000, attorney docket number 044463-0023;
- Application Serial No. 09/354,058, titled HIERARCHICAL BACKUP AND RETRIEVAL SYSTEM, filed July 15, 1999, attorney docket number 044463-0014;

- Application Serial No. 09/774,302, titled LOGICAL VIEW WITH GRANULAR ACCESS TO EXCHANGE DATA MANAGED BY A MODULAR DATA AND STORAGE MANAGEMENT SYSTEM, filed January 30, 2001, attorney docket number 044463-0040;
- 5 • Application Serial No.09/876,289, titled APPLICATION SPECIFIC ROLLBACK IN A COMPUTER SYSTEM, filed June 6, 2000, attorney docket number 044463-0029; and
- Application Serial No. 09/038,440, titled PIPELINED HIGH SPEED DATA TRANSFER MECHANISM, filed March 11, 1998, attorney docket number 10 4982/6;
- each of which applications is hereby incorporated herein by reference in this application.

BACKGROUND OF THE INVENTION

The invention disclosed herein relates generally to backup storage systems and methods for computer data. More particularly, the present invention relates to managing shadow copies of a volume.

The server operating system by Microsoft Corp. of Redmond, Washington called XP/.NET Server contains an integrated application for making shadow copies. Such shadow copies are also known as "snapshots" and can either be hardware or software copies depending on the snapshot program being used. Common snapshot programs include the previously-mentioned XP/.NET Server snapshot program by Microsoft, the TimeFinder snapshot program by EMC Corp. of Hopkinton, Massachusetts, and the EVM snapshot program by Compaq Computer Corp. of Houston, Texas.

Generally, when a shadow copy is taken, a new logical volume is exposed on 25 the machine that is an exact image of the original volume. While changes can continue to

occur on the original volume, the new volume is a static, point-in-time view of the data.

Since shadow copies persist on a user's workstation, a different network machine, etc. the shadow copies provide the ability to have multiple versions of data ready for recovery at a moment's notice. Minimal restore time, and the downtime associated therewith while the

- 5 restore operation is being performed, is therefore provided since there is no need to mount external media, such as tape or optical media, to stream data back therefrom.

Although shadow copying offers quick backup and recovery capability, the snapshots are stored on relatively expensive media, such as a fast hard drive, a redundant array of independent disks ("RAID") system. RAID refers to a set of two or more ordinary

- 10 hard disks and a specialized disk controller. The RAID system copies data across multiple drives, so more than one disk is reading and writing simultaneously. Fault tolerance is achieved by mirroring, which duplicates the data on two drives, and parity, which calculates the data in two drives and stores the results on a third. A failed drive can be swapped with a new one, and the RAID controller rebuilds the lost data on the failed drive. Some backup

- 15 storage systems copy backups to slower media, such as slow hard drives, tape drives, etc.; however, the downtime associated with a backup and recovery for such systems is increased.

Moreover, backup copies are formatted or compressed for optimum utilization of storage

media. Restoring backup copies further require the extra step of unformatting or

uncompressing the backup copy for use by the computer system. There is therefore a need for

- 20 a backup storage system which minimizes the downtime associated with a backup and restore operation while taking advantage of less expensive media.

Additionally, the software products available to create shadow copies, such as the XP/.NET, TimeFinder, etc., lack efficient management of shadowed copies. For instance, administrators in many instances must track shadowed copies, remember which original

- 25 volume corresponds to particular shadowed copies, what data existed on them, when a copy

operation occurred, if a copy should be destroyed, etc. There is therefore a need for methods, systems, and software products that enable efficient management of shadowed copies.

BRIEF SUMMARY OF THE INVENTION

The present invention provides methods, systems, and software products that

5 enable efficient creation, management, and recovery of shadowed copies and quick recovery volumes of primary volumes or applications. Particularly, the invention provides methods and systems for creating a quick recovery volume and snapshot images of primary volumes and application data from a single interface.

In one aspect of this invention, a computer readable medium which stores

10 program code is provided that when executed on a computer, causes the computer to perform a method for creating a quick recovery volume of a primary data set used by a first computer in a backup storage system. In one embodiment, the method includes identifying a snapshot image of the primary data set generated by a snapshot application, and creating the quick recovery volume of the primary data set from the snapshot image of the primary data set. The

15 method also including controlling transfer of data from the first computer to an archival storage unit. In one embodiment, the data set is a primary volume or application data. The quick recovery volume may also be a disk-to-disk data-block-level replication of the data set. The quick recovery volume may be an incremental backup of a previous quick recovery volume of the primary data set.

20 In one embodiment, the program code includes an agent module and a storage manager module. The agent module enables data transfer from the first computer to the archival storage unit and the storage agent module interfaces the agent module and the archival storage unit. The agent module may be an intelligent agent module, which enables data transfer of the primary data set for a specific application. The program code may further

25 provide a quick recovery agent that evokes a snapshot application to create the snapshot

image of the primary data set. The primary data set may include a plurality of primary volumes, at least one primary volume and at least one application data set, or a plurality of application data sets. The scope of the primary data set may be defined as a sub-client of the first computer. The details to create the quick volume may be provided in a quick recovery 5 policy data structure.

In one embodiment, the program code causes the first computer to automatically select a destination volume for the quick recovery volume of the primary data set from a pool of available volumes. The destination volume for the quick recovery volume of the primary data set may be selected based on storage space available on an available 10 volume in comparison to storage space needed for the quick recovery volume, the selected volume capacity exceeding that needed for the quick recovery volume of the primary data set and closer to a capacity needed than other available volumes.

In one aspect of this invention, a computer system is provided that includes an archival storage unit, and a programmed computer for controlling data transfer from the 15 computer to the archival storage unit to create a quick recovery volume of a primary data set. The computer may provide a snapshot image of the primary data set, and create the quick recovery volume of the primary data set from the snapshot image for the primary data set. The archival storage unit may be connected to the client computer over a communication network. The computer system may also include at least one server computer 20 communicatively coupled to the programmed computer and the archival storage unit. The server may be programmed for controlling data transfer from the computer to the archival storage unit.

In one aspect of this invention, a method for creating a quick recovery volume of a primary data set of a first computer, is provided that includes the steps of creating a 25 snapshot image of the primary data set and creating the quick recovery volume of the primary

data set from the snapshot image of the primary data set. The step of creating the quick recovery volume of the primary data may be creating the quick recovery volume as a disk-to-disk data-block-level replication of the primary data set. The quick recovery volume of the primary data set may also be an incremental backup of a previous quick recovery volume of the primary data set. The quick recovery volume may further be a block-level copy of the primary data set from the snapshot image of the primary data set.

In one embodiment, the method of creating a primary recovery volume includes the step of synchronizing with an operating system to flush all data of the primary data set to an archival storage unit during the creation of the snapshot image of the primary data set. Synchronizing may include suspending input to a disk containing the primary data set during the creation of the snapshot image of the primary data set. The method may further include resuming input to the disk containing the primary data set upon creation of the snapshot image of the primary data set. The steps of suspending and resuming may be accomplished automatically or manually with user-supplied command line commands during presnapshot and post-snapshot phases. In one embodiment, the snapshot images of the primary data set are also indexed. In one embodiment of the invention, the method further includes the step of deleting the snapshot image of the primary data set at a selected time. The selected time may be immediately after a copy phase or after a persistence period.

In one aspect of this invention, a computer readable medium storing programming code is provided. The programming when executed causes a computer to present a snapshot manager that interfaces with a backup storage system. The snapshot manager enables users to browse snapshot images and enables users to recover snapshot images stored in the backup storage system. The snapshot manager may interface with a volume snapshot service for creating a snapshot image. The stored program code may further include a snapshot requester program module that packages data for the snapshot image. The

packaged data may be communicated to the volume snapshot service for creating the snapshot image.

In one embodiment the programming code includes a snapshot requester program module and a snapshot writer program module, which may be directed by the 5 snapshot requester program module to package data for the snapshot image. The writer program module may package data for a specific application, which may then be communicated to the volume snapshot service for creating the snapshot image.

In one embodiment, the snapshot manager program module indexes snapshot images. The snapshot indexing enables copying, deleting, displaying, browsing, changing 10 properties, or recovering snapshot images. The snapshots may be indexed in terms of objects native to particular applications.

The snapshot management tool may integrate with existing backup systems, such as the Galaxy™ backup system provided by CommVault Systems of Oceanport, New Jersey and further described in Application Number 09/610,738. The present invention 15 leverages the indexing technology and ‘point-in-time’ browse and recovery capability of such systems to manage shadow copies. Alternatively, the snapshot management tool may act as a stand-alone management tool for basic snapshot management not requiring integration with existing backup systems, such as the CommVault Galaxy™ backup system and others.

20

BRIEF DESCRIPTION OF THE DRAWINGS

The invention is illustrated in the figures of the accompanying drawings which are meant to be exemplary and not limiting, in which like references are intended to refer to like or corresponding parts, and in which:

Fig. 1 is a block diagram depicting the software components and communication paths of program code stored on a computer readable medium for a backup storage system according to an embodiment of the invention;

Fig. 2 is a block diagram of a typical storage system model;

5 Fig. 3 is a computer system according to an embodiment of the invention;

Fig. 4 is a flow diagram of a method of creating a backup copy of a primary data set of a client computer according to an embodiment of the invention;

Fig. 5 is a block diagram depicting the software components and communication paths of program code stored on a computer readable medium for a backup storage system with snapshot capability according to an embodiment of the invention; and

10 Figs. 6 and 7 are browser style user interface screens according to an embodiment of the invention.

DETAILED DESCRIPTION

Referring to Fig. 1, software components of a computer readable medium for 15 use in creating quick recovery volumes of a primary data set of a client computer in a backup storage system 100, according to an embodiment of this invention, includes at least one agent 102, such as an intelligent data agent 104, e.g., the iDataAgent™ module available with the Galaxy™ system, a quick recovery agent 108, a media agent 106, e.g., the MediaAgent™ module available with the Galaxy™ system, etc., and at least one storage manager 110, e.g., 20 the CommServe StorageManager™ module also available with the Galaxy™ system. A primary data set generally denotes a volume, application data, or other data being actively used by of a client computer. A volume is generally an area of storage on a data storage device, which may be the whole storage area or portions thereof. An agent 102 generally refers to a program module that provides control and data transfer functionality to client 25 computers. A client generally refers to a computer with data, e.g., a primary data set, that

may be backed up, such as a personal computer, workstation, mainframe computer, host, etc. Intelligent data agent 104 refers to an agent for a specific application, such as Windows 2000 File System, Microsoft Exchange 2000 Database, etc., that provide control and data transfer functionality for the data of the specific applications. A plurality of agents 102, such as 5 intelligent data agents 104 or quick recovery agents 108, may be provided for and/or reside on each client computer, for example, where the client computer includes a plurality of applications and a file system or systems for which a quick recovery volume may be created.

A quick recovery volume as is used herein generally to denote a full replica of an original volume. A full replica implies an unaltered copy of the primary data set, such as 10 an unformatted or uncompressed copy, as is typically the case with backup copies. This enables faster recovery for a client computer by simply mounting or pointing to the quick recovery volume.

In one embodiment, where a quick recovery volume is being created for one or more volumes and/or applications on a client computer, a sub-client may be created. A sub-client generally refers to a defined set of parameters and policies that define the scope of the data set, such as the volumes or applications that are going to be copied, recovered, or otherwise managed. A sub-client generally contains a subset of the volumes and applications of the client. Multiple sub-clients may be created for a client computer and the sub-clients may overlap such that they include common data sets between them. A quick recovery 15 volume for a client, sub-client, or a plurality of sub-clients may be directed to point to a quick recovery policy or policies data structures, which provide the details for creating a quick recovery volume, such as how snapshots for the volumes or applications are created, copied, snapshot and quick recovery volume persistence, data pruning, the destination volume of the quick recovery volume, etc.

The destination volume for the quick recovery volume may be specified to be copied to specific volumes, or may be selected automatically from a pool of available volumes. The quick recovery agent 108 or the media agent 106, in one embodiment, selects an available volume as the destination volume where the quick recovery volume will be stored. The quick recovery agent 108 may select the volume at random or target a volume according to the storage space available on a particular volume in comparison to the space needed for the quick recovery volume. Once the volume is selected, it is removed from the pool of available volumes. This may be accomplished, for example, by the media agent 106 determining the capacity needed for a quick recovery volume, determining the capacity of the available volumes, and selecting the volume with a capacity exceeding that needed for the quick recovery volume and closer to the capacity needed than the other volumes.

A media agent 106 generally refers to a software module that provides control for archival storage units 112, such as tape library, a RAID system, etc., and facilitates local and remote data transfer to and from the archival storage units, or between the clients and the archival storage units. The media agent 106 may interface with one or more agents 102, such as the intelligent data agent 104 or quick recovery agent 108, to control the data being copied from a client computer, such as a primary volume or application data, to the storage volumes. A primary volume generally refers to a volume of a client computer that is the original source of the data, e.g. the primary data set, for the quick recovery volume. Data generally refers to information that may be stored on a storage device, including the file system, applications, and information related thereto. For example, the media agent 106 may interface with a quick recovery agent 108 to act as a copy manager 116, which manages the copying of data from primary volumes 114 to the quick recovery volumes 118. A storage manager 110 generally refers to a software module or application that interfaces the plurality of agents, clients, storage units, etc., and in one embodiment, coordinates and controls data flow between them.

The primary volumes 114 and the quick recovery volume 118 may be stored via a variety of storage devices, such as tape drives, hard drives, optical drives, etc. The storage devices may be local to the client, such as local drives, or remote to the client, such as remote drives on a storage area network (“SAN”) or local area network (“LAN”) environment, etc.

5 A quick recovery agent 108 generally refers to a software module that provides the ability to create quick recovery volumes 118. The quick recovery agent 108 evokes a snapshot mechanism or interfaces with a snapshot manager that provides for the creation of a snapshot image of the primary data set, such as of a primary volume or of application data. In one embodiment, the snapshot image of the primary data set is accessed to create a quick
10 recovery volume of the primary data set.

In one embodiment, the quick recovery agent 108 is a stand-alone application that adds to or interfaces with snapshot image programming, such as XP/.NET, TimeFinder, etc., that create snapshots or shadowed copies of the primary data set for the creation of the quick recovery volume 118 of the primary data set. Alternatively, or in addition, the quick
15 recovery volume 118 is a disk-to-disk data-block-level volume or application data replication of a client computer.

In one embodiment, the snapshot images of the primary data set are stored on fast media, such as a fast hard drive or RAID system and the quick recovery volume is stored on slow media, such as a hard drive or a tape library. The client computer may be a stand-
20 alone unit or connected to an archival storage unit in a storage area network (“SAN”) or local area network (“LAN”) environment.

In one embodiment, an initial quick recovery volume is created by capturing a snapshot image of the primary data set and creating the quick recovery volume from data stored on the primary volume. The quick recovery volume is subsequently updated to include
25 changes to the primary data set by referencing changes appearing in subsequent snapshot

images of the primary data set. This may be accomplished by tracking data changes between snapshot images. The snapshot images may include the changed data or simply track the data that has changed. The quick recovery volume may then be incrementally updated in accordance with the data changes or with reference to the tracked changes in the snapshot 5 images.

In one embodiment, the quick recovery volume of the primary data set is an incremental backup. That is, the operation for creating or updating the quick recovery volume is performed by incrementally copying, from the primary volume or application data to a previous a snapshot image or images of the primary data set, blocks of data that have 10 been modified since the pervious snapshot images. Alternatively, creating or updating a quick recovery volume is accomplished by incrementally copying data that has changed from a primary volume to the quick recovery volume with reference to changes tracked in the snapshot image. In one embodiment, a plurality of snapshot images of the primary date set are created and the data from the snapshots is incrementally stored between the snapshots to 15 provide redundant quick recovery.

Referring to Fig. 2, a typical storage system model 200 for a client computer includes a plurality of layers, such as an operating system layer 202, an applications layer 204, etc. The operating system layer further includes a plurality of layers or sub-layers, e.g., a physical disk layer 206, a logical volume manager ("LVM") layer 208, a file system layer 20 210, etc. The physical disk layer 206 denotes physical storage devices, such as a magnetic hard drive or disk array. The LVM layer 208 refers to logical disk volume management, which allows efficient and flexible use of the physical disk storage, for example, by permitting the physical disk to be divided into several partitions that may be used independently of each other. Some LVMs combine several physical disks into one virtual 25 disk. LVMs may write metadata, such as partition tables, to reserved areas of the physical

disk. The LVM virtual disks are transparently presented to upper layers of the system as block-addressable storage devices having the same characteristics as the underlying physical disks.

The file system layer 210 represents a higher-level logical view of the data,
5 which typically consists of a hierarchy of nested directories, folders, and files, and metadata. The features and attributes of files may vary according to the particular file system in use. For example, an NTFS, i.e., a Windows NT file system, tracks ownership and per-user access rights on each file, whereas FAT, i.e., file allocation table, file systems do not provide security features. Above the file system, and outside the scope of the operating system in
10 general, is the applications layer 204. The applications layer 204 includes application software, such as a word processor program, etc. which interface with the file system provided by the operating system to store data. Sophisticated software, such as database management systems (“DBMS”), may use special file system features or even raw logical volumes, and employ measure to protect the consistency of data and metadata. The
15 consistency of the data and metadata may be maintained during the creation of a quick recovery volume with writers particular to specific applications or file systems. Writers are described in more detail below.

Referring to Fig. 3 a computer system, according to one embodiment of the invention, includes a client computer 302, such as a personal computer, a workstation, a
20 server computer, a host computer, etc. In one embodiment, the client computer 302 contains programming which enables the creation of local quick recovery volumes of a primary volume or of application data. That is, the client computer 302 deploying the programming creates a quick recovery volume or volumes that are stored or copied locally at the client computer, such as on a local hard drive, tape drive, optical drive, etc. In one embodiment, the
25 programming is deployed on at least one client computer 302 connected over a

communications network 304, such as a LAN or SAN, to at least one archival storage unit 112, such as a tape library, a stand alone drive 306, a RAID cabinet, etc. In one embodiment, the client computer 302 include programming, such as an agent 102 or a storage manager 110, that provides data transfer functionality from the client computer 302 to the archival storage unit 112. In one embodiment, at least one of the client computers 302 also acts as a server computer 304. The server computer 302 generally contains programming, such as a media agent 106 or a storage manager 110 to control data transfer between the client computers 302 and the archival storage units 112. In one embodiment, at least two client computers 302 act as server computers 304; at least one server providing media agent functionality and at least one server providing storage manager functionality.

Referring to Fig. 4, a method for creating a quick recovery volume of a primary data set of a client computer 400, according to one embodiment of this invention, is performed in a plurality of phases. In one embodiment, the quick recovery volume 118 is created in two phases, a snapshot phase 408 and a copy phase 414. Each of the snapshot and copy phases may include a plurality of accompanying phases. For instance, the snapshot phase 408 may include a presnapshot 406 and a post snapshot phase 410. Similarly, the copy phase 404 may include a pre-copy phase 412 and a post copy phase 416.

In the snapshot phase 402, the quick recovery agent 108 synchronizes with the applications, if any, and the operating system to ensure that all data of the primary data set to be backed up is flushed to the archival storage unit or destination disk, where the quick recovery volume of the primary data set will be stored, and to ensure that the primary disk where the primary data set, such as the primary volume or application data, is located is not modified during the creation of the snapshot image, step 422. This may be accomplished for instance by suspending input or output to the primary disk containing the primary data set, step 420, which will ensure that the file system and metadata remain unchanged during the

copy operation. In one embodiment, the quick recovery agent evokes the snapshot mechanism to create a snapshot image of the primary data set. The snapshot mechanism or snapshot manager may be a software module, an external snapshot application, such as CommVault Software Snapshot, XP/.NET, TimeFinder, etc., or a combination thereof. Once 5 the snapshot image of the primary data set is created, application access to primary disk may resume, step 424, and update data on the primary disk as necessary, while the copy operation for the quick recovery volume 118 is in progress or is pending. In one embodiment, once the snapshot image is made the snapshot image is indexed step 425. Indexing generally denotes associating information with a snapshot image that may be useful in managing snapshot 10 image, such as the date the snapshot image was created, the lifespan of the snapshot, etc.

During the copy phase, the quick recovery volume is created from the snapshot image of the primary data set so that any suspension in the input or output to primary disk may be minimized. This may be accomplished by the media agent 106 referring or pointing 15 the quick recovery agent 108 to the snapshot volume or copy as the source of the data for the quick recovery volume 118. The relevant agent or agents may then package the data from the snapshot volume or copy, communicate the packaged data to the media agent 106 or quick recovery agent 108, and the media agent 106 or quick recovery agent 108 may send the data to the quick recovery volume 118 for copy. Packaging generally denotes parsing data and logically addressing the data that is to be used to facilitate the creation of the quick recovery 20 volume. For example, where a snapshot of the Microsoft Exchange application is to be created, the Exchange specific intelligent agent will parse the relevant data from the primary disk or disks containing the application data and logically address the parsed data to facilitate rebuilding the parsed data for the quick recovery volume. In one embodiment, the copy phase is performed after a specified amount of time has lapsed, such as a day, two days, etc., or at a 25 specified time. In yet another embodiment, a plurality of snapshot images of the primary data

set may be created at various times and the oldest snapshot image is copied to the quick recovery volume.

- The method of creating quick recovery volumes 118 may differ for particular applications. For example, for the Microsoft Exchange 2000 application, prior to suspending
5 input/output to the storage group associated with the application, the entire storage group is dismounted automatically during the snapshot phase 402 and remounted automatically when the snapshot is ready. For the SQL 2000 database, the database may be frozen automatically and released when the snapshot is ready. In one embodiment, suspend and resume functions for particular applications may be accomplished with user-supplied command line commands
10 or script, which may be entered during the presnapshot phase 406 or post-snapshot phase 410. Command line commands or script may further be entered to perform any additional processing that may be required, such as steps to synchronize with an application not supported by the quick recovery agent, or where an alternate host backup is desired, a command to mount the given volume onto the alternate host can be specified.
15 In one embodiment, during the copy phase 414, the quick recovery agent 108 performs a block-level copy of the primary data set from the snapshot image to the destination disk or volume, step 426, which becomes the quick recovery volume 118. Command line commands or script may also be provided during a precopy phase 412 and a post-copy phase
416.
20 Users may recover data from a snapshot image or the quick recovery volume 118, step 428. In one embodiment, recovery generally entails suspending input or output to the disk containing the quick recovery volume of the primary data set where the data will be copied from, step 430, restoring the primary data set to the primary volume , step 432, and resuming input or output to the disk, step 434. Restoring the primary data set, such as
25 application data, includes mounting a volume containing the snapshot image of the primary

data set, such as the primary volume or application data, or mounting a quick recovery volume 118 of the primary data set in place of the primary volume, or replacing individual files, folder, objects, etc. to the primary volume from the quick recovery volume. In one embodiment, where a backup copy of the primary data set replaces a primary volume, input or 5 output to disk is not suspended. In one embodiment, the method of creating a quick recovery volume 118 includes an unsnap phase 418, which generally entails deleting the snapshot image that was created during the snapshot phase 408. The snapshot may be deleted at a specified time, such as immediately after the creation of the quick recovery volume or after a persistence period, which period so that the resources may be available for future quick 10 recovery volume creations.

In one embodiment, backup software, such as the quick recovery agent 108 interfaces with a snapshot manager to access a snapshot image of the primary data set for the creation of a quick recovery volume 118. A snapshot manager may be a stand-alone application or program module that controls the creation and management of snapshot images 15 of primary volumes or of application data. Referring to Fig. 5, a snapshot manager 503, according to one embodiment of this invention, is a program module, such as a snapshot manager agent, which interfaces with the backup programming, such as the quick recovery agent 108. The snapshot manager may be an intelligent agent in that it manages snapshot for a specific application, e.g., Windows 2000 File system, Exchange, Oracle, etc., a plurality of 20 which can be installed on any client computer to create snapshot copies of a plurality of applications' data.

The creation and management of a snapshot image of the primary data set may be further accomplished with a snapshot requestor 502. In one embodiment, the snapshot requester 502 is a program module that generally packages data of particular applications or 25 of primary volumes for. In one embodiment, when the creation of a snapshot image is

requested by the quick recover agent 508, for example, the snapshot requestor 502 communicates with a snapshot writer 504 and directs the writers to package the data requested for the snapshot image. In one embodiment, snapshot writers 504 are application specific modules designed to package data from individual applications, such as Windows 5 2000 file system, Microsoft Exchange, Oracle, etc. After the snapshot writer 502 packages the data, the data is communicated to a volume snapshot service 505, which actually creates the snapshot image of the primary data set. In one embodiment, the snapshot writer 504 communicates the packaged data to the snapshot requestor 502, which then passes the packaged data to the volume snapshot service 505.

10 The volume snapshot service 505 is either a software snapshot application from a software snapshot provider, such as Microsoft .NET Server, or a hardware snapshot application from a hardware snapshot provider, such as EMC or Compaq. The software snapshot image applications will, in one embodiment, create a space efficient copy that is exposed as a separate logical volume using a copy-on-write technique. Hardware snapshot 15 image applications, typically accompanied with a RAID cabinet, create a mirror or clone copy of application data or primary volumes. Once the volume snapshot service 505 has taken the snapshot image, the snapshot data is passed to the snapshot manager 503, which indexes the snapshot image enabling snapshot management. Indexing generally denotes associating snapshots with information that may be useful in managing snapshots, such as the date the 20 snapshot was created, the lifespan of the snapshot, etc. Managing generally includes, but is not limited to, copying, deleting, displaying, browsing, changing properties, or restoring the snapshots or data therein. Indexing generally provides point-in-time browse and management, such as recovery, capability of the snapshot images and of the quick recovery 25 volumes. Users can choose to persist or retain snapshot images well beyond the lifetime of the requesting application or module. The snapshot manager 503 may then communicate the

snapshot data to the quick recovery agent 108 for copying to the quick recovery volume 118, or to the media agent 106 for copying to the archival storage unit 112.

In one embodiment, the present invention implements a high-performance data mover for performing a disk-to-disk data transfer. Data mover may also perform server-less 5 data transfer using extended copy to create secondary or auxiliary copies over the communication network, e.g., SAN or LAN. In one embodiment, an extended copy command acts as a copy manager, which is embedded on a SAN component, such as a gateway, router, tape library, etc. Alternatively, the copy manager is a program module that interfaces with the backup storage system. In another embodiment, hardware snapshots are 10 mounted on an alternate host to perform a server-free backup. This effectively allows a user to convert a software snapshot image to the equivalent of a hardware snapshot image that can be persisted or retained.

In one embodiment, aware technology, described in United States Patent Application Number 09/610,738, is incorporated into snapshot image programming or 15 volume snapshot services to make the applications aware. In other words, the intelligent data agent 104 makes objects that are native to particular applications part of the snapshot image, which enables the user to perform actions in terms of the applications' objects. This enables, for instance, browsing snapshot images of volumes consisting of Exchange data that will be visible in terms of storage groups and stores, rather than just a volume consisting of 20 directories and files. For particular applications, such as Exchange or SQL Server, a further level of detail with regard to the objects may be included, such as paths to Exchange objects, such as Storage Groups or stores, or paths to SQL objects, such as databases, file-groups, or files. This information may be used at the time of browsing to determine if any of the existing snapshot volumes contain copies of the objects of interest so that they may be 25 presented to the user for recovery. In other embodiments similar application-aware

configurations are provided for applications such as Lotus Notes, Oracle, Sharepoint Server, etc.

- In one embodiment, the snapshot manager 503 is accessible to a user with an appropriate user interface screen which enables the creation and management of snapshot images or quick recovery volumes of a primary volume or application data, contained on a client computer. Actions that are available to users include (1) create a snapshot image, e.g., snap, at a specified time (2) snap and persist for a period, (3) specify the destination volume of a snapshot image and where the image should persist (for software snapshot), (4) specify or change the period a snapshot image should persist, (5) browse existing snapshot images, (6)
- 10 recover a snapshot image to a specified volume, and (7) destroy or delete a snapshot image.
- Browsing generally denotes enabling a user to view information for particular snapshots. For example, browsing enables a user to view the available snapshots for a particular volume or application data and information related thereto. Recovering generally refers to replacing the primary data set with data from a snapshot image or quick recovery volume, which includes
- 15 mounting a volume containing the snapshot image or quick recovery in place of the primary volume, replacing application data on the primary disk from a snapshot or quick recovery volume, etc. In one embodiment, during a restoration, data may be retrieved from a plurality of quick recovery volumes, snapshot images, or a combination thereof. For example, data may be retrieved from a snapshot image and a quick recovery volume.
- 20 In one embodiment, a quick recovery volume or snapshot image of the primary data set may be used on a permanent basis as the primary data set, e.g., the primary volume. For example, a user may choose to run an application, such as Exchange, from the quick recovery volume permanently and future backup operations for the application will reflect the quick recovery volume as the primary volume. Setting up the backup operations as the
- 25 replacement for the primary volume may be accomplished by identifying a quick recovery

- policy for the backup operation and the backup volumes available to the client, and releasing a volume from the pool of available volumes. This method of recovering a primary volume or application provides a faster method of recovering data since the data transfer from backup copies to the primary copy is effectively eliminated. Moreover, recovering from a quick
- 5 recovery volume is a faster alternative than traditional backup techniques since the quick recovery volume does not have to be unformatted or uncompressed in order for the client computer to use the data. This method may be performed manually or automatically, and relevant tables or databases, such as the snapshot table may be amended to reflect the replacement volume as the primary volume or application for future backup operations.
- 10 In one embodiment, users may drill down to view particular folder, files, etc., or to view particular objects native to applications. In one embodiment, users are able to specify, with regard to a quick recovery copy, (1) whether the snapshot image should persist after the quick recovery volume, (2) if the image should persist, for how long, and (3) the location of the persistent storage for the image. In yet another embodiment, users are able to
- 15 (1) request a snapshot image and a quick recovery volume, just a snapshot image, or just a quick recovery volume, (2) request a software snapshot image and optionally specify that it be converted to a hardware snapshot image, (3) request that the hardware snapshot image persist for a certain period of time, (4) recover data from a snapshot image at the volume level, e.g., the whole volume, or sub-volume level, e.g., individual folders, files, objects, etc., and (5)
- 20 make another copy of a snapshot image on the SAN.

In one embodiment, snapshot information that has been indexed or associated with snapshot images by the snapshot manager 503, is tracked in at least one table or database, e.g., snapshot table, which is accessible to the backup storage system 100 or the storage manager 110. The snapshot table, in one embodiment, contains information for every

25 volume or copy that has been configured for every client, application, or, sub-client,

indicating the snapshot images that are currently available for a particular volume, application, sub-client, etc. The snapshot information preferably includes a timestamp that indicates when a snapshot was created and a time interval that indicates how long the snapshot should persist. The snapshot table may be accessed by any one of the program
5 modules for managing and controlling the quick recovery volumes.

In one embodiment, an application or module, such as the snapshot manager, enables the following functionality. When a snapshot image of a volume or application is being performed, the application suspends input or output to a disk, determines which applications reside on the primary volume, engages the relevant writers for the particular
10 applications, performs or directs a snapshot, packages the snapshot data, and resumes the input/output to the disk. The application also makes appropriate entries into the snapshot table for the given client. When a backup is being performed, the application identifies the content, identifies the volumes involved, identifies the applications involved, engages all the writers involved, and performs or directs the snapshot, and performs or directs the copying to
15 the quick recovery volume 118. Appropriate entries are made into the snapshot table for the given client and volumes involved. In some embodiments of the present invention, as part of the creation of a snapshot, the application enters into the snapshot table parameters relating to how long the snapshot should persist. When a snapshot is destroyed or deleted, the application performs or directs the deletion of the snapshot and updated the snapshot table
20 accordingly.

If the snapshot is a software snap and persistent storage has been identified to convert it to a hardware snap, the application will first perform a fast copy of the data, such as with DataPipe and backup APIs, to accomplish the data movement, and then update the snapshot table. Every time the application is evoked, it re-discovers the volumes on the given
25 client and ensures that any new volumes are added to the default sub-client of an agent, such

as the snapshot manager intelligent agent. The application can also be called as part of a recovery operation. In which case a copy of the data is made from one volume to another. In addition, if operating system data, such as metadata, is involved, certain writers may have to be engaged to ensure a correct data restoration.

5 In one embodiment, pruning of data is also enabled, such as snapshot images that have expired or their persistence period lapsed. Pruning may be scheduled to run periodically, such as weekly, monthly, etc. If snapshot images are present and their persistence period has lapsed, the snapshot image is destroyed or deleted and the snapshot table is updated accordingly.

10 While the discussion above assumes that the application or module, such as the snapshot manager, encapsulates logic to manipulate the built-in shadow copy mechanism in Windows .NET Server, the application can easily encapsulate the same logic for any hardware snapshot which will be recognized by those skilled in the art. The application may be implemented in conjunction with plug-in modules, dynamic link libraries ("DLLs"), that
15 will each support a different snapshot program such as .NET, TimeFinder, EVM, etc.

One embodiment of the present invention provides a user interface screen for users to browse and recover data, such as from snapshot images, quick recovery volumes, primary copies, backup copies, etc., as of a point-in-time. Browsing and recovery may be client, sub-client, volume, and application specific, and may be at the volume level or at the
20 sub-volume level. Volume level recovery refers to replication of entire volumes, whereas sub-volume level refers to recovery at a folder, file, or object level. Referring to Fig. 6, a browser interface screen 600, according to one embodiment of this invention, includes a plurality of frames, such as directory frame 602 and a contents frame 604. The directory frame generally provides a list of all available drives, partitions, volumes, snapshots, backups,
25 etc. and the file folders therein, of a client computer in a hierarchical arrangement. The

contents frame 604 generally lists the contents of any item appearing in the directory frame 602, such as folders, files, or objects. The contents may be displayed by highlighting any one of the items in the directory frame 602. By selecting the “My Snapshots” folder, for example, the contents of the snapshots folder 612 are displayed in the contents frame 604. The 5 contents may be displayed with relevant details, such as the date of creation, persistence, association, the capacity of the volume, etc. In one embodiment, the user may change the properties of a snapshot, such as how long a particular snapshot will persist, the location, etc., and the user may direct the creation of another volume or copy of a software snapshot using, for example, CommVault data movers.

10 In one embodiment, users may specify a point-in-time for which browsing and restoration may occur. In that instance, the browser application determines if there are any existing snapshot images present as of the point-in-time specified. Snapshot data found to be available as of the point-in-time the user specified is displayed to the user. Snapshot data is displayed if it exists and qualifies as valid data as of the point-in-time. If the browser 15 application does not find a snapshot, backup copies, such as primary copies, and secondary copies, and quick recovery volumes are presented or accessed for data recovery or restoration. If the user chooses to drill down a given snapshot image, quick recovery volume, or backup copy to see the contents therein, such as by selecting or double-clicking an item, the item is displayed at the requesting client computer in an appropriate user interface screen, such as in 20 an interface screen provided by the application associated with the item. Association generally refers to the relationship between a file and the application that created it.

 The snapshot folder and contents displayed at the user interface as of a certain point-in-time may be provided by browse logic that will check the snapshot table to see if there is a snapshot available as of that point-in-time for volumes or copies of interest. If there 25 is a snapshot available, data relating to the content of the snapshot is displayed accordingly.

Application specific objects are mapped to data files or directories and this mapping is stored in database tables. This allows for an application-specific view of objects on the snapshot when the snapshot is browsed or recovered. For example, where a snapshot of C:\ volume has been created, browsing under heading “My Snapshots” may reveal a C:\ volume that is a 5 snapshot image of the C:\ volume. Alternatively, snapshot images may be designated with different labels. For example, the snapshot volume of C:\ may be V:\ with a label indicating that V:\ is a snapshot volume of C:\. Drilling down through the snapshot and the folders therein may reveal the file, folders, or objects, which may be viewed, recovered, restored, deleted, etc. For example, a file “important.doc” appearing in the snapshot of the C:\ volume 10 may be viewed with a document viewer, deleted, recovered, or restored to the primary volume.

Referring to Fig. 7, a browser interface screen for browsing snapshot images according to an embodiment of the present invention displays a particular client 702 as a folder, for example, a folder for the client squid.commvault.com. At least one subfolder may 15 be displayed showing the application or applications available with respect to the client for creating backup copies, browsing, and recovery. For example, the “Exchange 2000 Database,” the “File System,” and “SQL Server 2000.” Selecting one of the subfolders, such as “Exchange 2000 Database” reveals subfolders therein, such as a “SnapShot Data” subfolder 706, which provides snapshot data for the selected application. Further drilling 20 down through the subfolders will cause to be displayed in a hierarchical layout the snapshots available, such as “SnapShot 1,” application objects 710, such as the “Information Store,” “First Storage Group”, “Mailbox Store”, “Public Folder Store”, etc. A user may then perform a single click recovery or restore, or creation of a snapshot image or quick recovery volume of a primary volume or application data, or of any backup copy by selecting the level from 25 which the data displayed there under will be backed up or recovered. For example, by

selecting the “First Storage Group” and right clicking, the user will be presented with an activity window 712, which allows the user to select the “Recover” function. In this instance, by selecting “Recover” all data related to the objects appearing under the “First Storage Group” will be recovered. The user may choose to recover a single file, or a single object,

- 5 such as a single Exchange store or SQL Server database. The restoration may be implemented with a fast data mover, such as CommVault’s DataPipe™, described in detail in Application No. 09/038,440, which will move data from disk to disk. Additionally a disk-to-disk server-less data mover can be implemented as well.

The user interface screen may be used to recover a data set, such as the
10 primary volume or application data, from a given snapshot copy or quick recovery volume. If an entire data set is to be recovered, the snapshot image or quick recovery volumemay be substituted for the original volume, such as with the recovery process described above. Recovery of items smaller than a volume can be accomplished by using traditional file copy techniques or with Windows Explorer, such as by copying and pasting the desired files or
15 objects. Depending on the files or objects being recovered, the recovery process may involve identifying which writers were involved at the time of the snapshot and then engaging them to accomplish the restoration.

The user interface screen may also be used to request destruction of a given
snapshot copy on a single item bases or automatically upon the lapse of the persistence
20 period. Basic information for each snapshot copy destroyed, or otherwise, may be stored for purposes of tracking and display. This information, in one embodiment, is stored in an MSDE database, but can also be stored in other similar data structures.

Some of the embodiments of the present invention leverage existing features
of the CommVault Galaxy backup system. It will be recognized by those skilled in the art,
25 however, that the embodiments of the present invention may be applied independently of the

Galaxy system. While the invention has been described and illustrated in connection with preferred embodiments, many variations and modifications as will be evident to those skilled in this art may be made without departing from the spirit and scope of the invention, and the invention is thus not to be limited to the precise details of methodology or construction set forth above as such variations and modification are intended to be included within the scope of the invention.

What is claimed is:

1. A computer readable medium storing program code which when executed on a computer, causes the computer to perform a method for creating a quick recovery volume of a primary data set used by a first computer in a backup storage system, the method comprising:
 - 5 identifying a snapshot image of the primary data set generated by a snapshot application,
creating the quick recovery volume of the primary data set from the snapshot image of the primary data set; and
controlling transfer of data from the first computer to an archival storage unit.
 - 10 2. The computer readable medium of claim 1, wherein the data set comprises at least one of a primary volume and application data.
 3. The computer readable medium of claim 1, wherein the quick recovery volume comprises a disk-to-disk data-block-level replication of the data set.
 4. The computer readable medium of claim 3, wherein the quick recovery volume 15 comprises an incremental backup of a previous quick recovery volume of the primary data set.
 5. The computer readable medium of claim 1, wherein the program code comprises an agent module and a storage manager module, wherein the agent module enables data transfer from the first computer to the archival storage unit and the storage agent module 20 interfaces the agent module and the archival storage unit.
 6. The computer readable medium of claim 5, wherein the agent module comprises an intelligent agent module, which enables data transfer of the primary data set for a specific application.

7. The computer readable medium of claim 1, wherein the program code comprises a quick recovery agent that evokes a snapshot application to create the snapshot image of the primary data set.

8. The computer readable medium of claim 1, wherein the primary data set comprises one of: a plurality of primary volumes, at least one primary volume and at least one application data set, and a plurality of application data sets; a scope of the primary data set defined as a sub-client of the first computer.

9. The computer readable medium of claim 1, wherein the computer readable medium stores a quick recovery policy data structure providing details to create the quick recovery volume of the primary data set.

10. The computer readable medium of claim 1, wherein the program code causes the first computer to automatically select a destination volume for the quick recovery volume of the primary data set from a pool of available volumes.

11. The computer readable medium of claim 10, wherein the destination volume for the quick recovery policy of the primary data set is selected based on storage space available on an available volume in comparison to storage space needed for the quick recovery volume of the primary data set, the selected volume capacity exceeding that needed for the quick recovery volume of the primary data set and closer to a capacity needed than other available volumes.

12. A computer system comprising:
an archival storage unit; and
a programmed computer for controlling data transfer from the computer to the archival storage unit to create a quick recovery volume of a primary data set, the computer providing a snapshot image of the primary data set, and creating the quick recovery volume of the primary data set from the snapshot image for the primary data set.

13. The computer system of claim 12, the archival storage unit connected to the programmed computer over a communication network.

14. The computer system of claim 13, comprising at least one server computer communicatively coupled to the programmed computer and archival storage unit and programmed for controlling data transfer from the computer to the archival storage unit.

15. A method for creating a quick recovery volume of a primary data set of a first computer, the method comprising:

creating a snapshot image of the primary data set; and
creating the quick recovery volume of the primary data set from the snapshot image of the primary data set.

16. The method of claim 15, comprising creating the quick recovery volume of the primary data set as a disk-to-disk data-block-level replication of the primary data set.

17. The method of claim 15, comprising creating the quick recovery volume of the primary data set as an incremental backup of a previous quick recovery volume of the primary data set.

18. The method of claim 15, comprising synchronizing with an operating system to flush all data of the primary data set to an archival storage unit during the creation of the snapshot image of the primary data set.

19. The method of claim 18, wherein synchronizing comprises suspending input to a disk containing the primary data set during the creation of the snapshot image of the primary data set.

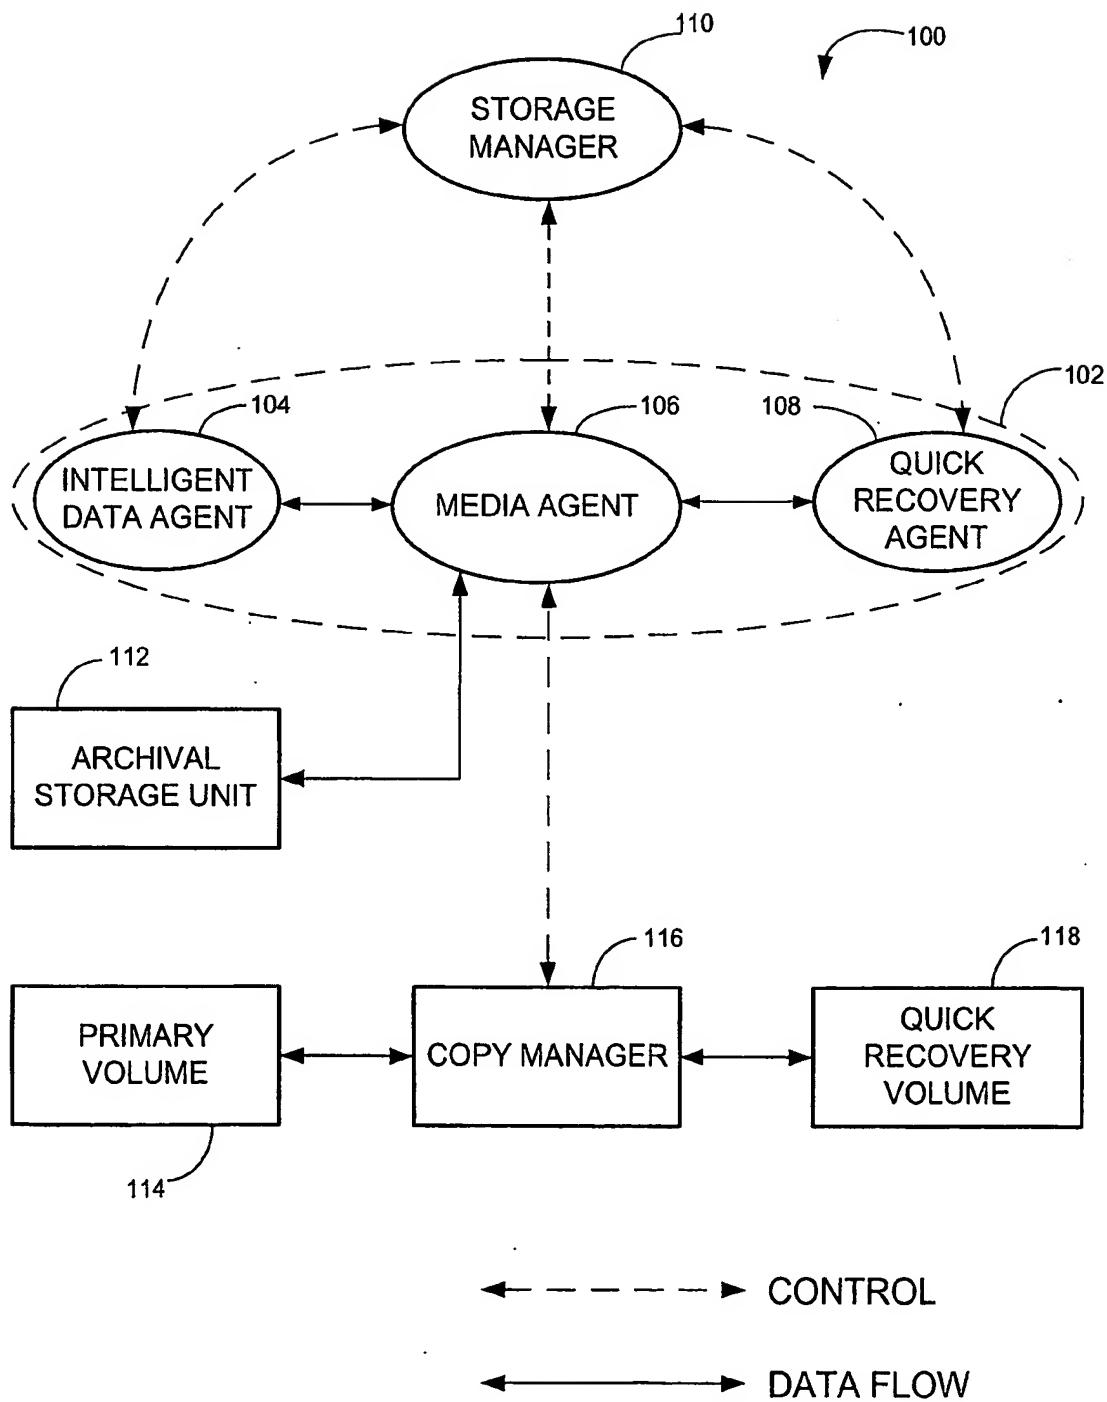
20. The method of claim 19, comprising resuming input to the disk containing the primary data set upon creation of the snapshot image of the primary data set.

21. The method of claim 20, wherein suspending and resuming are accomplished one of automatically, and manually with user-supplied command line commands during presnapshot and post-snapshot phases.
22. The method of claim 16, comprising performing a block-level copy of the primary data set from the snapshot image of the primary data set.
5
23. The method of claim 16, comprising indexing the snapshot image of the primary data set.
24. The method of claim 16, comprising deleting the snapshot image of the primary data set at a selected time, wherein the selected time comprises one of immediately
10 after a copy phase and after a persistence period.
25. A computer readable medium storing programming code, which when executed causes a computer to present a snapshot manager that interfaces with a backup storage system, the snapshot manager enabling users to browse snapshot images and enabling users to recover snapshot images stored in the backup storage system.
15
26. The computer readable medium of claim 25, the snapshot manager interfacing with a volume snapshot service for creating a snapshot image.
20
27. The computer readable medium of claim 26, wherein the stored program code further comprises a snapshot requester program module that packages data for the snapshot image, the packaged data communicated to the volume snapshot service for creating the snapshot image.
25
28. The computer readable medium of claim 26, comprising a snapshot requester program module, and a snapshot writer program module directed by the snapshot requester program module to package data for the snapshot image, the writer program module packaging data for a specific application, the packaged data communicated to the volume snapshot service for creating the snapshot image.

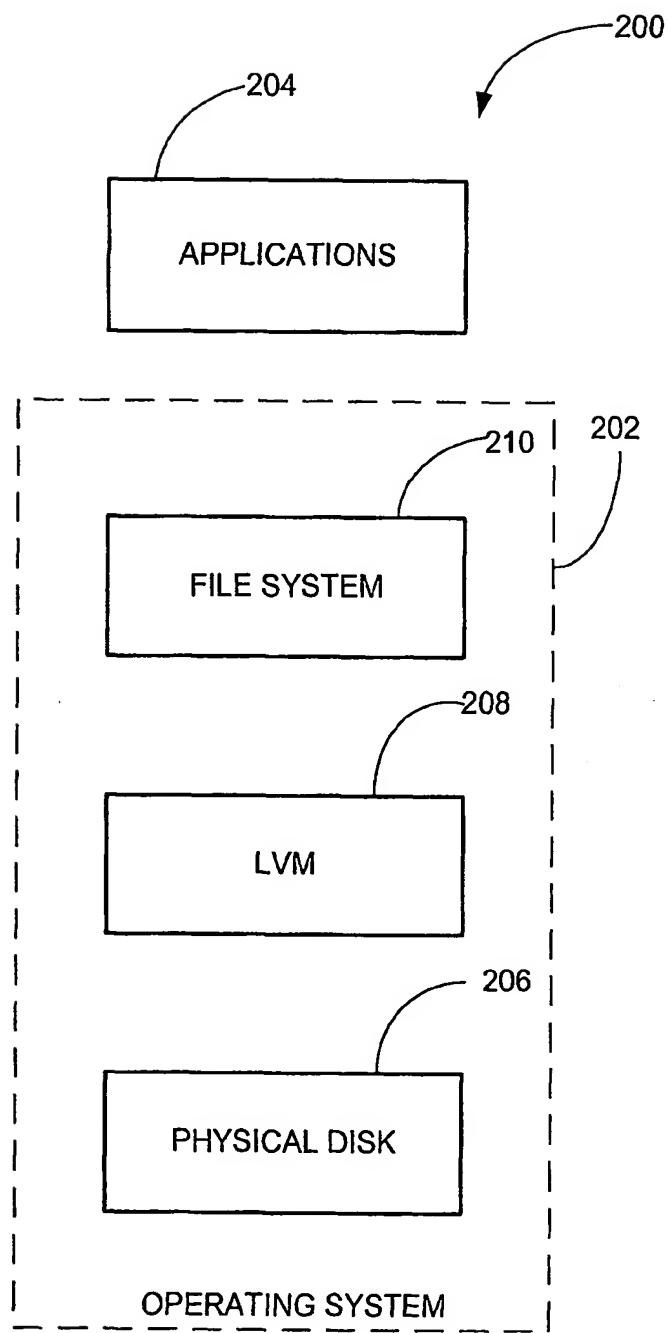
29. The computer readable medium of claim 26, wherein the snapshot manager program module indexes snapshot images, snapshot indexing enabling at least one of copying, deleting, displaying, browsing, changing properties, and recovering snapshot images.

30. The computer readable medium of claim 29, wherein the indexing comprises
5 indexing in terms of objects native to particular applications.

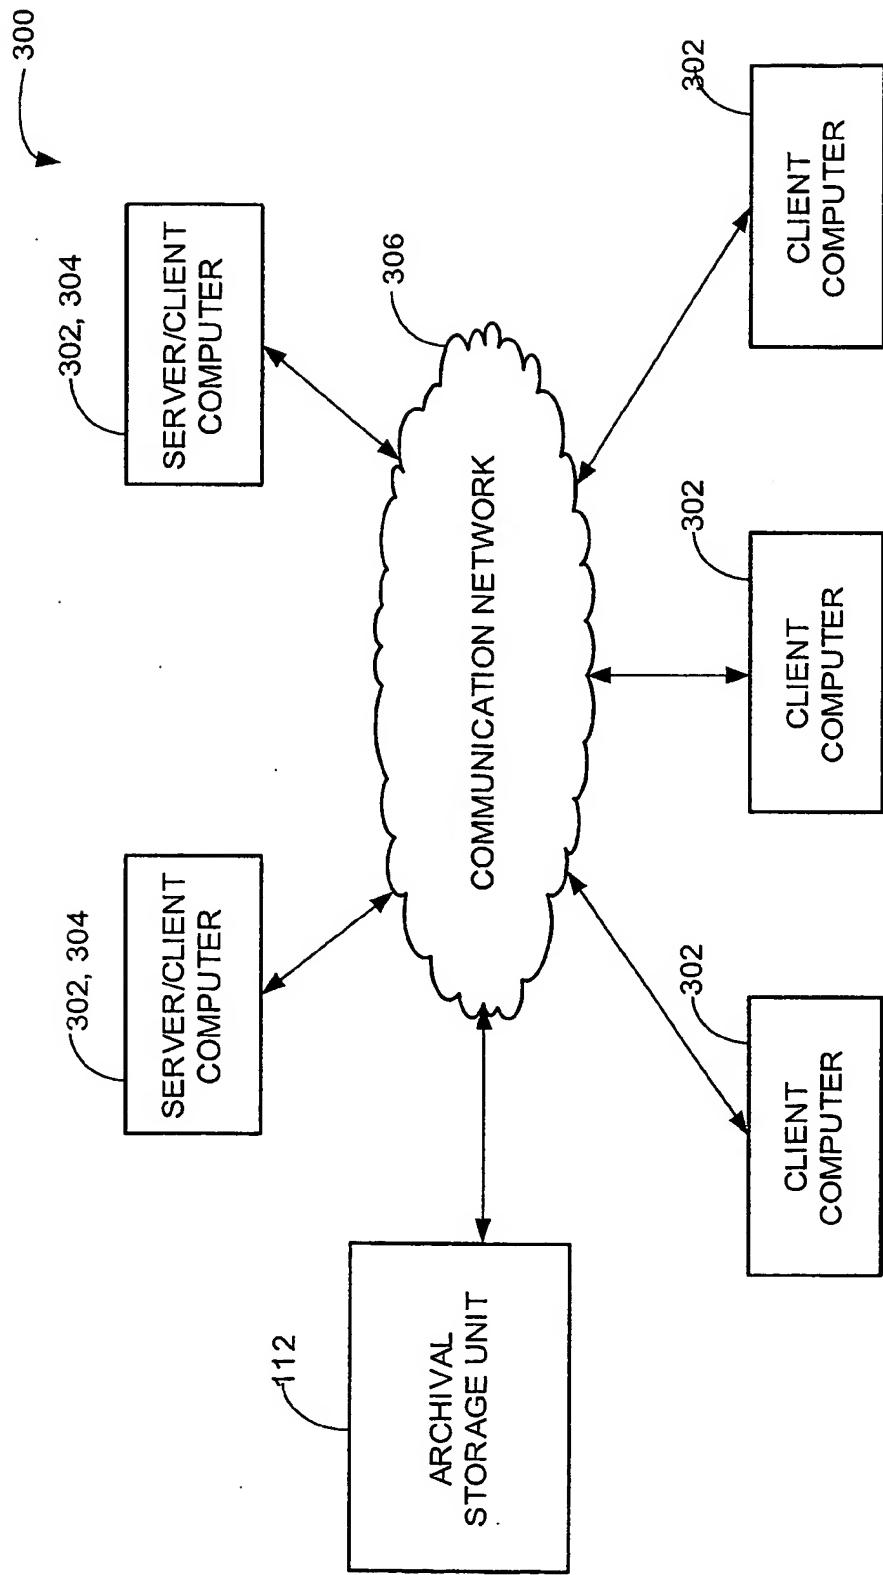
1/7

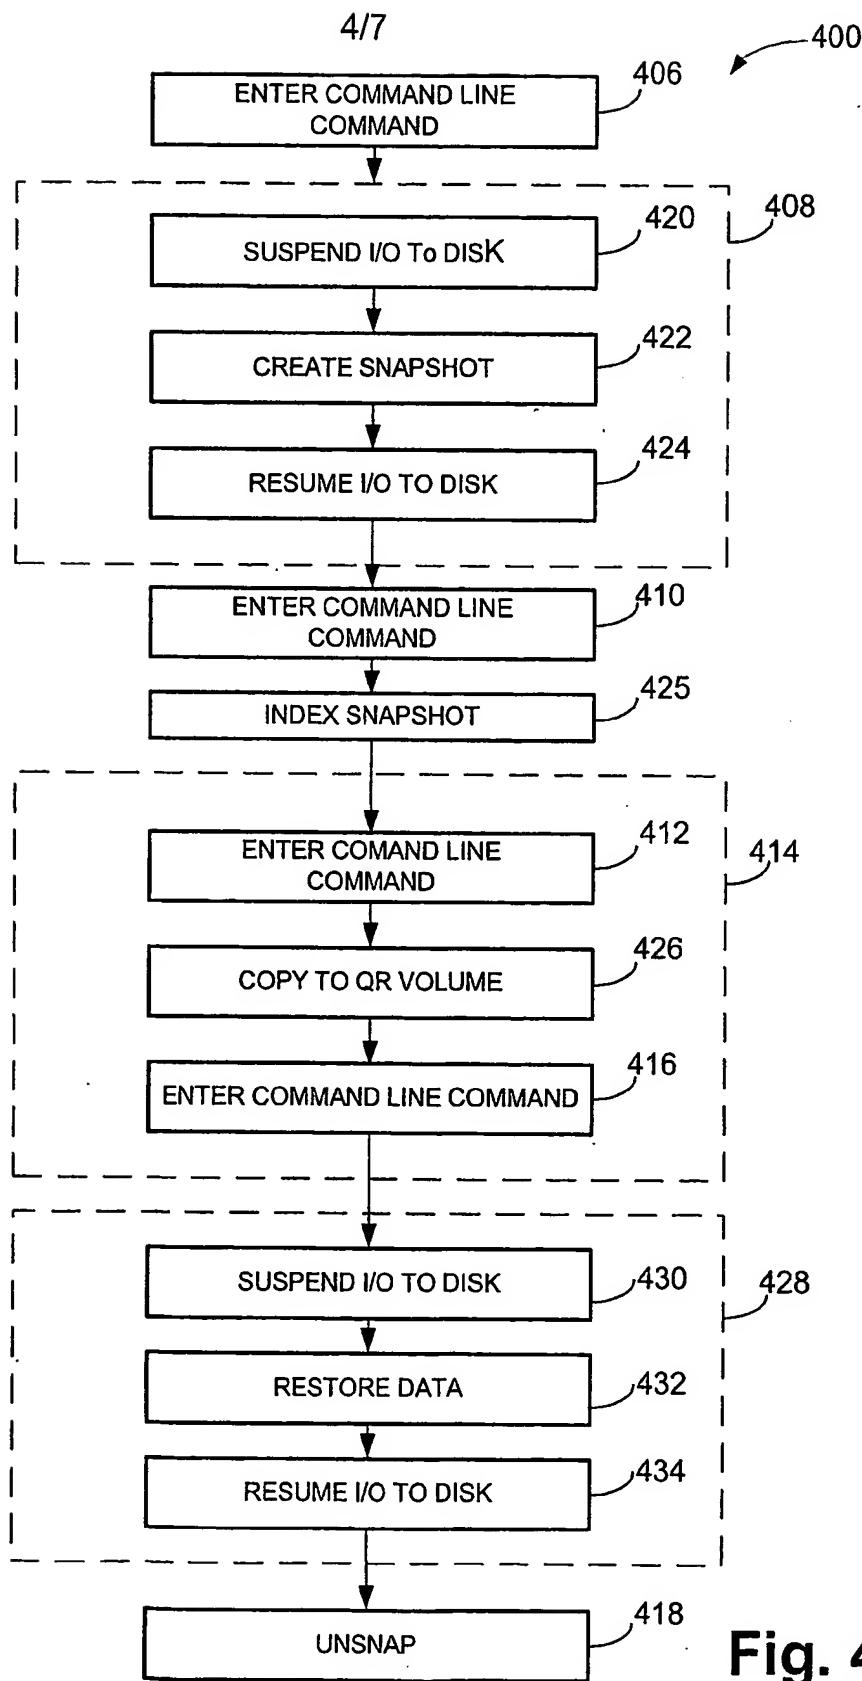
**Fig. 1**

2/7

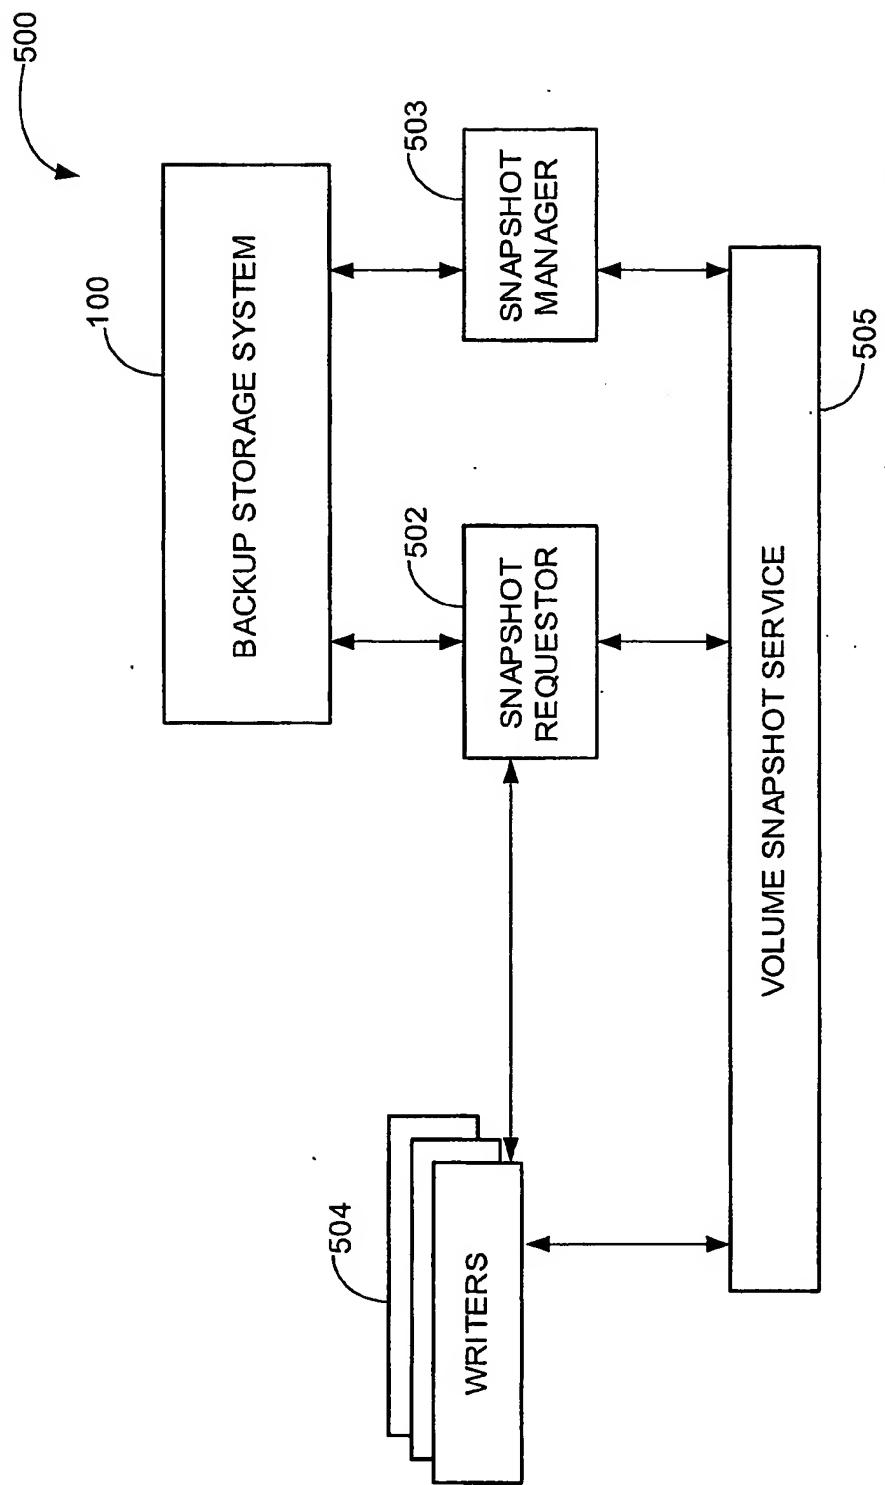
**Fig. 2**

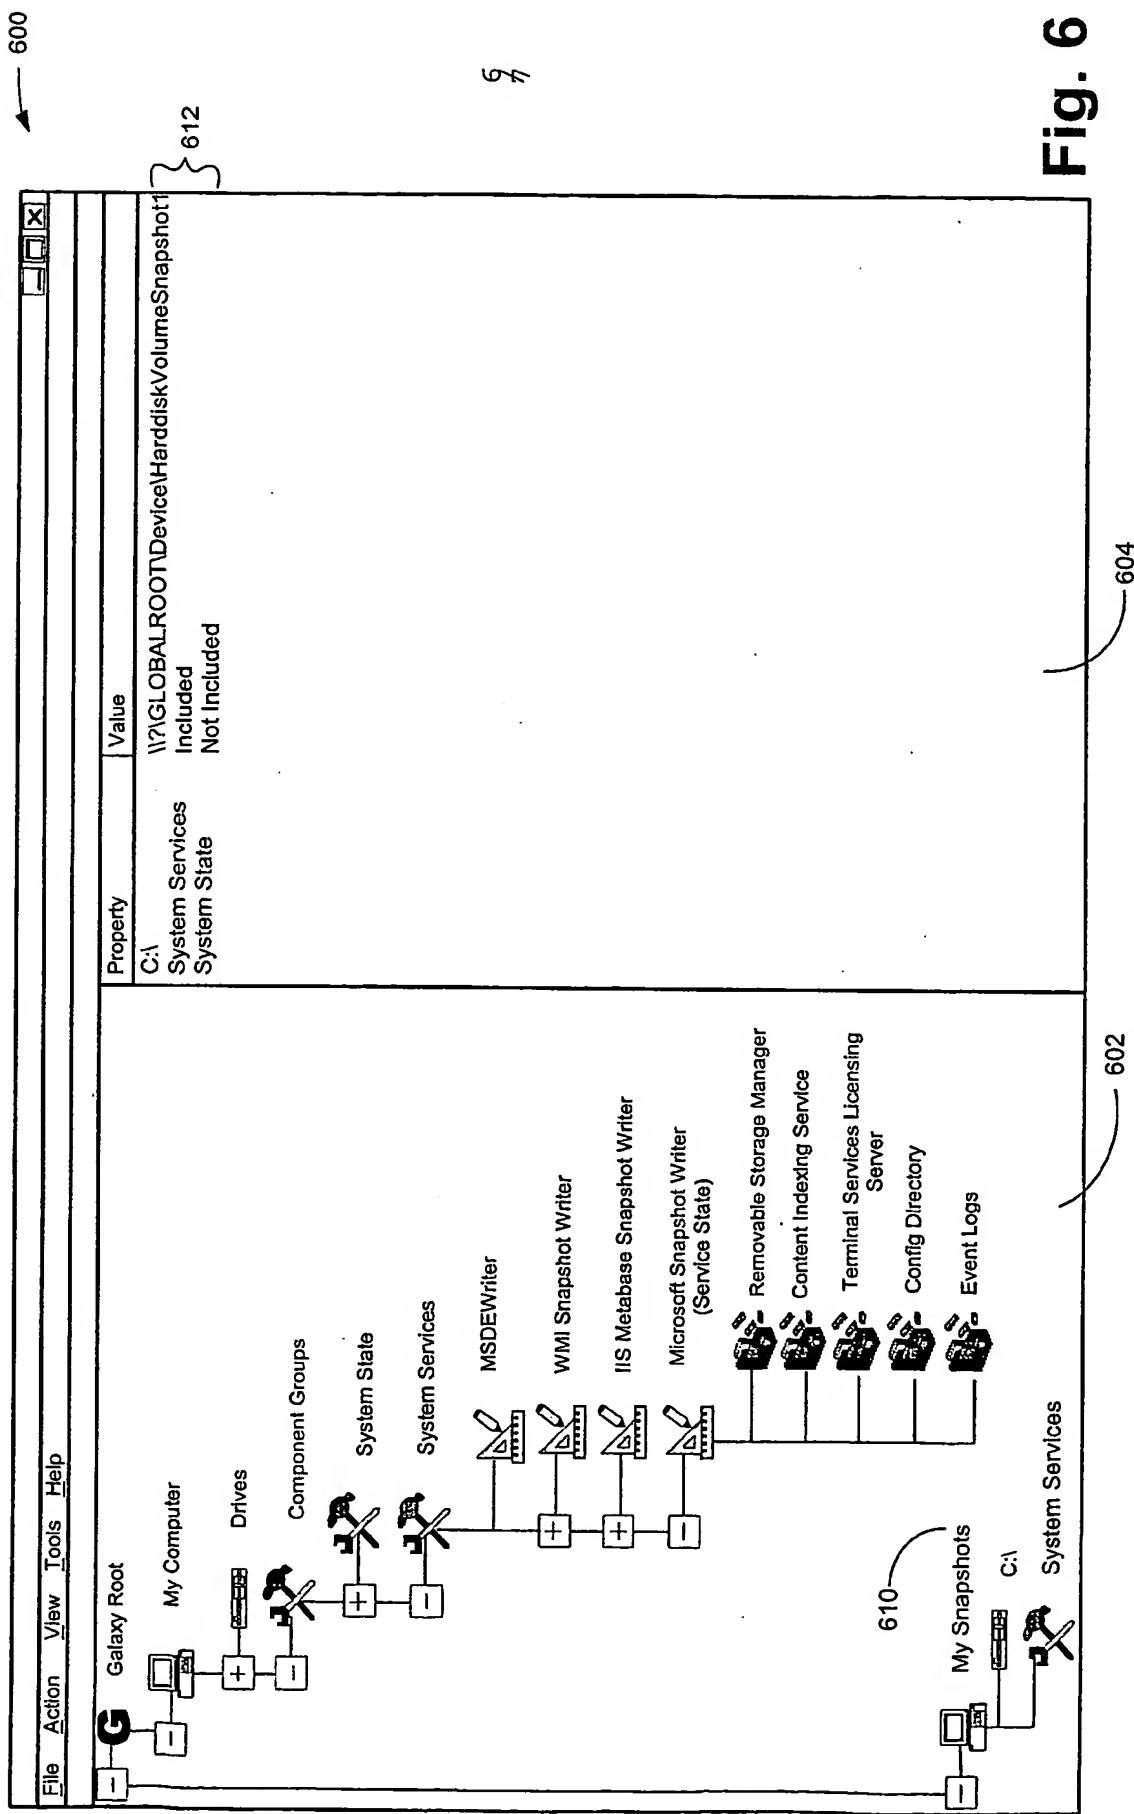
3/7

**Fig. 3**

**Fig. 4**

5/7

**Fig. 5**



7/7

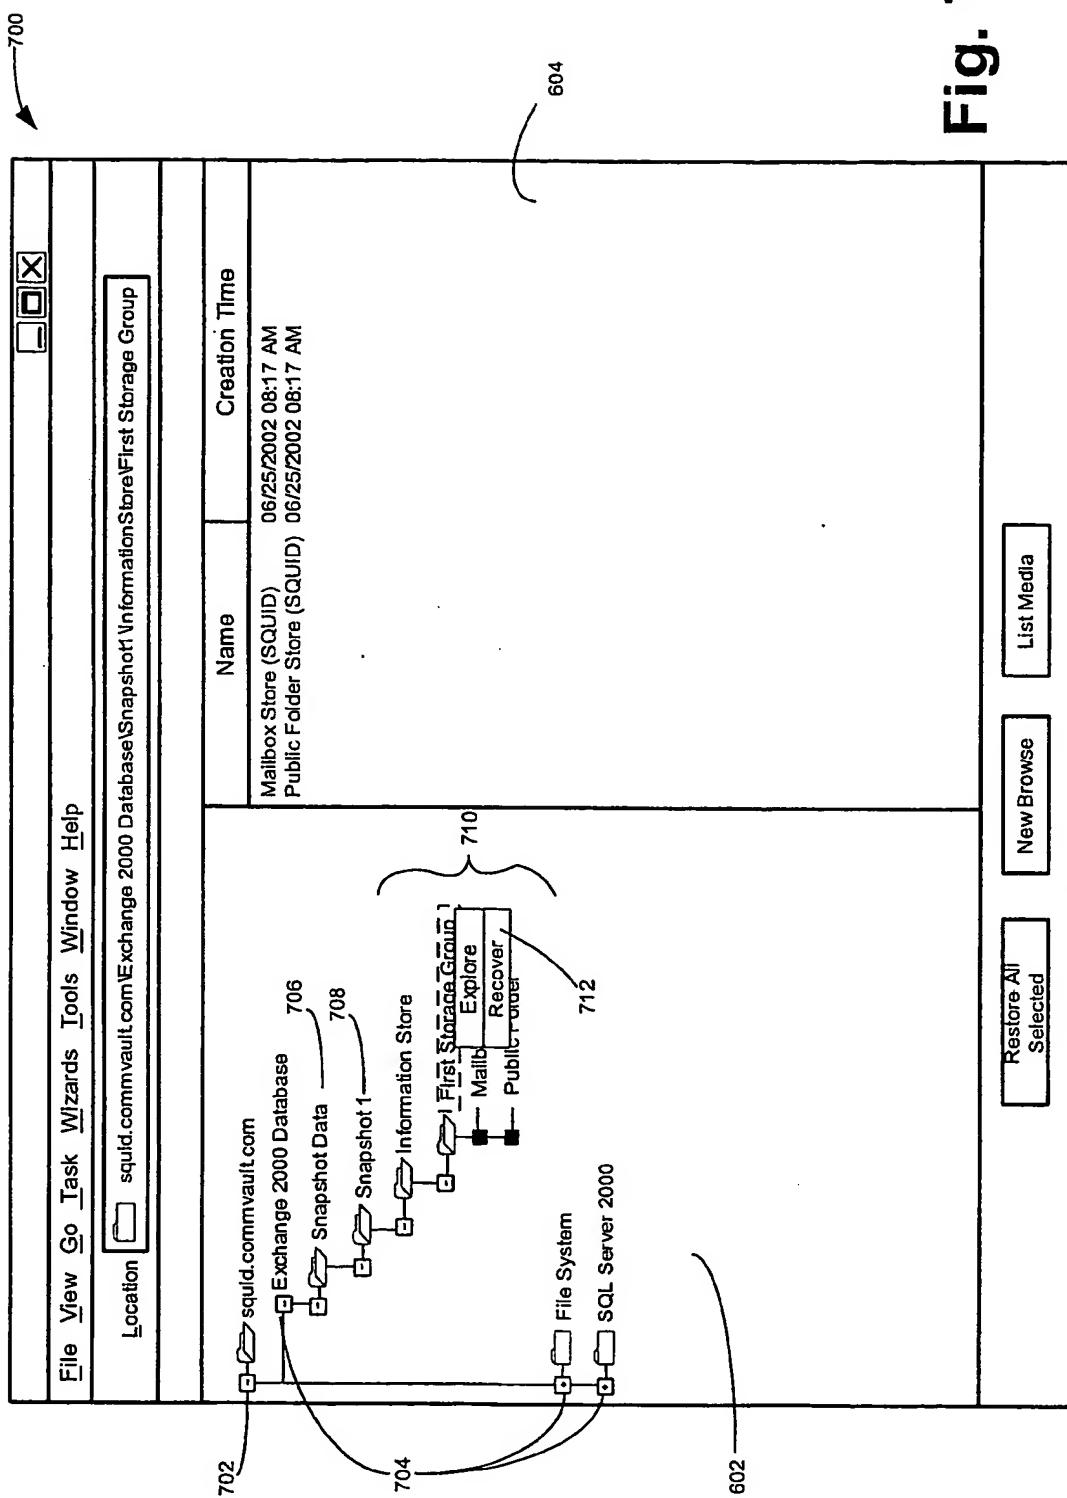


Fig. 7

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US02/31206

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : H02H 3/05; G06F 12/00
 US CL : 714/6; 711/161,162; 707/204

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 U.S. : 714/6; 711/161,162; 707/204

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 EAST

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6,374,363 B1 (Wu et al.) 16 April 2002 (16.4.2002), see column 3 lines 20-25; column 11 lines 13-25	1-2, 4-9, 13, 16, 18-21
---		-----
Y		3, 10, 17, 22-30
Y	US 6,487,561 B1 (Ofek et al.) 26 November 2002 (26.11.2002), column 3 lines 10-20; column 24 lines 39-41	3, 10, 17, 22
X	US 6,434,681 B1 (Armangau) 13 August 2002 (13.8.2002), column 3 lines 1-15; column 7 lines 17-25; column 13 lines 18-25 and 55-67; column 14 lines 1-60	1-2, 7-11, 13-16, 19-21, 23-24
---		-----
Y		3, 17, 22, 25-30
Y	US 6,072,490 A (Bates et al.) 06 June 2000 (6.6.2000), column 2 lines 10-20	25-30

Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier application or patent published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&"	document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means		
P document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search 21 January 2003 (21.01.2003)	Date of mailing of the international search report 26 FEB 2003
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231	Authorized officer Robert W. Beausoliel <i>for James R. Matthews</i>
Facsimile No.	Telephone No. (703) 306-5631